arXiv:1510.07758v1 [cs.DS] 27 Oct 2015

Fast Compatibility Testing for Rooted Phylogenetic Trees*

Yun Deng† David Fernández-Baca‡

Abstract

We consider the following basic problem in phylogenetic tree construction. Let $\mathcal{P} = \{T_1, \ldots, T_k\}$ be a collection of rooted phylogenetic trees over various subsets of a set of species. The tree compatibility problem asks whether there is a tree T with the following property: for each $i \in \{1, \ldots, k\}$, T_i can be obtained from the restriction of T to the species set of T_i by contracting zero or more edges. If such a tree T exists, we say that $\mathcal{P}$ is compatible.

We give a $\tilde{O}(M_{\mathcal{P}})$ algorithm for the tree compatibility problem, where $M_{\mathcal{P}}$ is the total number of nodes and edges in $\mathcal{P}$. Unlike previous algorithms for this problem, the running time of our method does not depend on the degrees of the nodes in the input trees. Thus, it is equally fast on highly resolved and highly unresolved trees.

1 Introduction

Building a phylogenetic tree that encompasses all living species is one of the central challenges of computational biology. Two obstacles to achieving this goal are lack of data and conflict among the data that is available. The data shortage is tied to the vast disparity in the amount of information at our disposal for different families of species and the limited amount of comparable data across families [15]. One approach to overcoming this obstacle begins by identifying subsets of species for which enough data is available, and building phylogenies for each subset. The resulting trees are then synthesized into a single phylogeny —a supertree— for the combined set of species. This approach, proposed in the early 90s [2, 14], has been used successfully to build large-scale phylogenies (see, e.g., [3, 10]).

Any attempt at synthesizing phylogenetic information from multiple input trees must deal with the potential for conflict among these trees. Conflict may arise due to errors, or due to phenomena such as gene duplication and loss, and horizontal gene transfer. A fundamental question is whether conflict exists at all; that is, does there exist a supertree that exhibits the evolutionary relationships implicit in each input tree? We can formalize this question as follows. Let $\mathcal{P} = \{T_1, \ldots, T_k\}$ be a collection of rooted phylogenetic trees, where, for each $i \in \{1, \ldots, k\}$, T_i is a phylogenetic tree for a set of species $L(T_i)$. The *tree compatibility problem* asks whether there exists a phylogenetic supertree T for the set of species $\bigcup_{i=1}^{k} L(T_i)$ such that, for each $i \in \{1, \ldots, k\}$, T_i can be obtained from $T|L(T_i)$ — the minimal subtree of T spanning $L(T_i)$ — by zero or more contractions of internal edges (that is, $T|L(T_i)$ is homeomorphic to T). If the answer is "yes", then $\mathcal{P}$ is said to be *compatible*; otherwise, $\mathcal{P}$ is *incompatible*.

Here we present an algorithm that solves the compatibility problem for rooted trees in $O(M_{\mathcal{P}} \log^2 M_{\mathcal{P}})$ time, where $M_{\mathcal{P}}$ is the total number of vertices and edges in the trees in $\mathcal{P}$. This running time is independent of the degrees of the internal nodes of the input trees.

*Supported in part by the National Science Foundation under grants CCF-1017189 and CCF-1422134.

†Department of Computer Science, Iowa State University, Ames, IA 50011, USA, yundeng@iastate.edu

‡Department of Computer Science, Iowa State University, Ames, IA 50011, USA, fernande@iastate.edu

Previous work. Aho et al. [1] gave the first polynomial-time algorithm for the rooted tree compatibility problem. Their motivation was not phylogenetics, but relational databases. Steel [17] was perhaps the first to note the relevance of Aho et al.'s algorithm to supertree construction. His version of the Aho et al. algorithm, which he called the BUILD algorithm, has been a major influence in later work, including the present paper.

Henzinger et al. [9] showed that one can check the compatibility of a collection $\mathcal{R}$ of rooted triples — that is, phylogenetic trees on three species — in $O(|\mathcal{R}| \log^2 |\mathcal{R}|)$ time. (The time bound stated in [9] is higher, but can be improved using a faster dynamic graph connectivity data structure [11].) Any collection of trees $\mathcal{P}$ can be encoded by a collection of rooted triples $\mathcal{R}(\mathcal{P})$, obtained by enumerating the restriction of each input tree to every three-element subset of its species set (see Section 2). If n denotes the total number of distinct species in $\mathcal{P}$, then we get a trivial upper bound of $|\mathcal{R}(P)| = O(n^3 k)$. We can improve on this by finding a *minimal* set $\mathcal{R}^*$ of rooted triples that define the input trees. If the trees are binary — *fully resolved*, in the language of phylogenetics —, then $O(n)$ triples suffice for each tree, giving us $|\mathcal{R}^*| = O(nk)$. If input trees admit non-binary — that is, *unresolved* — nodes, however, the number of triples needed per input tree is roughly proportional to n^2 (the precise bound depends on the sum of the products of the degrees of internal nodes and the degrees of their children [8]), giving us $|\mathcal{R}^*| = O(n^2 k)$. Of course, the extra step of finding $\mathcal{R}^*$ adds to the complexity of the algorithm.

The tree compatibility problem is related to the *incomplete directed perfect phylogeny problem* (IDPP). Indeed, any collection of k phylogenetic trees on n distinct species can be encoded as a problem of testing the compatibility of a collection of $O(M_{\mathcal{P}})$ "directed partial characters" on n species[1]. Intuitively, each such character encodes the species in the subtree rooted at some node in an input tree. There is a $\tilde{O}(nm)$ algorithm to test the compatibility of m incomplete characters [13], which can be adapted to yield a $\tilde{O}(nM_{\mathcal{P}})$ algorithm for tree compatibility.

When the input trees are unrooted, the tree compatibility problem becomes NP-hard [17]. Nevertheless, the decision version is polynomial-time solvable if k is fixed [4]; that is, the problem is fixed-parameter tractable in k. The proof of fixed-parameter tractability in [4] relies on Courcelle's Theorem [6], and thus is an existence proof, rather than a practical algorithm.

Our contributions. At a high level, our algorithm resembles BUILD [17, 16]. There are, however, important differences. BUILD relies on the *triplet graph*, whose nodes are the species and where there is an edge between two species if they are involved in a triplet (see Section 2). Our algorithm relies instead on intersection graphs of sets of species associated with certain nodes of the input trees. Our graphs allow a more compact representation of the triplets induced by the trees in $\mathcal{P}$ (see Section 3). The key to the correctness of our approach is the intimate relationship between the triplet graph and our intersection graph (see Lemma 5 of Section 3). We remark that intersection graphs have a long history of use in testing compatibility, beginning with the work of Buneman [5].

We also take ideas from other sources. From Pe'er et al.'s IDPP algorithm [13], we adapt the idea of a *semi-universal node*. Although the graphs used to solve IDPP and rooted compatibility are different, semi-universal nodes play similar roles in each case: they capture the notion of sets of nodes in the input trees that map to the same node in a supertree, if a supertree exists. The relationship between our algorithm and Pe'er et al.'s goes deeper. Our approach can be viewed as an algorithm for IDPP that takes advantage of the fact that our particular set of incomplete characters arises from a collection of trees.

Intersection graphs are a convenient tool to prove the correctness for our algorithm. They are less convenient for an implementation, because it is hard to maintain them dynamically, in the way our algorithm requires. The difficulty lies in recomputing set intersections whenever the graphs are updated. We avoid this

[1] For a precise definition of partial characters and IDPP, we refer the reader to Pe'er et al. [13].

by using *display graphs*, an idea that we borrow from the proof of the fixed-parameter tractability of unrooted compatibility [4]. The display graph of a collection $\mathcal{P}$ is obtained by identifying leaves in the input trees that have the same label. Display graphs provide all the connectivity information we need for our intersection graphs (see Lemma 6 of Section 4), but are easier to maintain.

Through our techniques, we achieve what, to our knowledge, is the first algorithm for rooted compatibility to achieve near-linear time under all input conditions, regardless of the degrees of the nodes in the input trees. This is an essential quality for dealing with large datasets.

Contents. Section 2 reviews basic concepts in phylogenetics, defines compatibility formally, and introduces triplets and the triplet graph. Section 3 presents our intersection graph approach to testing tree compatibility. Section 4 describes the implementation details needed to achieve the $O(M_{\mathcal{P}} \log^2 M_{\mathcal{P}})$ time bound. Section 5 contains some final remarks.

2 Preliminaries

For each positive integer r, $[r]$ denotes the set $\{1, \ldots, r\}$.

Phylogenetic trees. Let T be a rooted tree. We use $V(T)$, $E(T)$, and $r(T)$ to denote the nodes, edges, and the root of T, respectively. For each $x \in V(T)$, we use $\mathrm{Ch}(x)$ and $T(x)$ to denote the set of children of x and the subtree of T rooted at x, respectively. Suppose $u, v \in V(T)$. Then, u is a *descendant* of v if v lies on the path from u to $r(T)$ in T. Note that v is a descendant of itself. T is *binary*, or *fully resolved*, if each of its internal nodes has two children.

A (rooted) *phylogenetic tree* is a rooted tree T where every internal node has at least two children, along with a bijection λ that maps each leaf of T to an element of a set of *species*, denoted by $L(T)$. For each $x \in V(T)$, $L(x)$ denotes the set of species mapped to the leaves of $T(x)$; that is, $L(x) = \{\lambda(v) : v \text{ is a leaf in } T(x)\}$. $L(x)$ is called the *cluster* at x. Note that $L(r(T)) = L(T)$. The set of all clusters in T is $\mathrm{Cl}(T) = \{L(x) : x \in V(T)\}$.

The following lemma, adapted from [16, p. 52], is part of the folklore of phylogenetics.

Lemma 1. *Let $\mathcal{H}$ be a collection of non-empty subsets of a set of species X that includes all singleton subsets of X as well as X itself. If there exists a phylogenetic tree T such that $\mathrm{Cl}(T) = \mathcal{H}$, then, up to isomorphism, T is unique.*

Let T be a phylogenetic tree and A be a set of species. The *restriction* of T to A, denoted $T|A$ is the phylogenetic tree with species set A where $\mathrm{Cl}(T|A) = \{C \cap A : C \in \mathrm{Cl}(T) \text{ and } C \cap A \neq \emptyset\}$. Let T' be a phylogenetic tree. T *displays* T' if $\mathrm{Cl}(T') \subseteq \mathrm{Cl}(T|L(T'))$. Equivalently, T displays T' if $T|L(T')$ is homeomorphic to T'.

A *rooted triple* is a binary phylogenetic tree on three leaves. A rooted triple with leaves a, b, and c is denoted $ab|c$ if the path from a to b does not intersect the path from c to the root. We treat $ab|c$ and $ba|c$ as equivalent.

When restricted to the three-element subsets of its species set, a phylogenetic tree T induces a set $\mathcal{R}(T)$ of rooted triples, defined as $\mathcal{R}(T) = \{T|X : X \subseteq L(T), |X| = 3 \text{ and } T|X \text{ is binary}\}$.

Lemma 2 ([16, p. 119])**.** *Let T and T' be two phylogenetic trees. Then T displays T' if and only if $\mathcal{R}(T') \subseteq \mathcal{R}(T)$.*

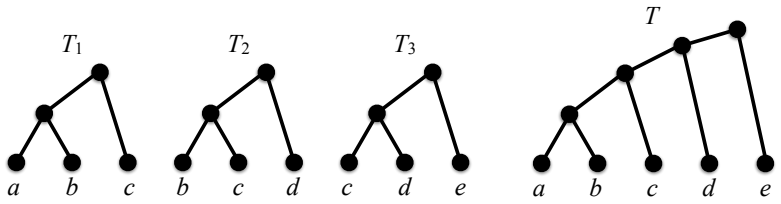

Figure 1: A profile $\mathcal{P} = \{T_1, T_2, T_3\}$ and a tree T that displays $\mathcal{P}$.

Profiles and compatibility. Throughout the rest of this paper $\mathcal{P} = \{T_1, \ldots, T_k\}$ denotes a set where, for each $i \in [k]$, T_i is a phylogenetic tree. We refer to $\mathcal{P}$ as a *profile*, and write $L(\mathcal{P})$ to denote $\bigcup_{i\in[k]} L(T_i)$, the *species set* of $\mathcal{P}$. We write $V(\mathcal{P})$ for $\bigcup_{i\in[k]} V(T_i)$, $E(\mathcal{P})$ for $\bigcup_{i\in[k]} E(T_i)$, and $\mathcal{R}(\mathcal{P})$ for $\bigcup_{i\in[k]} \mathcal{R}(T_i)$. Given a subset A of $L(P)$, $\mathcal{P}|A$ denotes the profile $\{T_1|A, \ldots, T_k|A\}$. The *size* of $\mathcal{P}$ is $M_\mathcal{P} = |V(\mathcal{P})| + |E(\mathcal{P})|$. Note that $M_\mathcal{P} = O(nk)$.

Profile $\mathcal{P}$ is *compatible* if there exists a phylogenetic X-tree T such that, for each $i \in [k]$, T displays T_i. If such a tree T exists, we say that T *displays* $\mathcal{P}$. See Figure 1.

The triplet graph. The *triplet graph* of a profile $\mathcal{P}$, denoted $\Gamma(\mathcal{P})$, is the graph whose vertex set is $L(P)$ and where there is an edge between species a and b if and only if there exists a $c \in L(P)$ such that $ab|c \in \mathcal{R}(\mathcal{P})$. The following observation concerning singleton profiles will be useful.

Observation 1. *Let T be a phylogenetic tree with $|L(T)| > 2$. Let $u_1, \ldots, u_p$ be the children of $r(T)$. Then, the connected components of $\Gamma(\{T\})$ are $L(u_1), \ldots, L(u_p)$, where $p \geq 2$.*

3 Testing Compatibility

Here we describe our compatibility algorithm and prove its correctness. We begin with some definitions.

Let U be a subset of $V(\mathcal{P})$ and let $L(U)$ denote $\bigcup_{u\in U} L(u)$. Then, $G_\mathcal{P}(U)$ denotes the graph with vertex set U and where $u, v \in U$ are joined by an edge if and only if $L(u) \cap L(v) \neq \emptyset$. That is, $G_\mathcal{P}(U)$ is the intersection graph of the clusters associated with the nodes in U. For each $i \in [k]$, let $U(i) = U \cap V(T_i)$. We say that U is *valid* if, for each $i \in [k]$,

(V1) if $|U(i)| \geq 2$, then there exists a node $v \in V(T_i)$ such that $U(i) \subseteq \mathrm{Ch}(v)$ and
(V2) $L(U(i)) = L(T_i) \cap L(U)$.

Observe that the set U_{init} defined as follows is valid.

$$U_{\mathrm{init}} = \{r(T_i) : i \in [k]\} \tag{1}$$

Note that $L(U_{\mathrm{init}}) = L(\mathcal{P})$. From this point forward, we assume that $G_\mathcal{P}(U_{\mathrm{init}})$ is connected. No generality is lost by doing so. To see why, observe that If $G_\mathcal{P}(U_{\mathrm{init}})$ is not connected, then $\mathcal{P}$ can be partitioned into a collection of species-disjoint profiles $\mathcal{P}_1, \ldots, \mathcal{P}_r$ such that $\mathcal{P}$ is compatible if and only if $\mathcal{P}_j$ is compatible for all $j \in [r]$.

The next observation follows from the definition of a valid set.

Observation 2. *If U is a valid subset of $V(\mathcal{P})$, then, for each $i \in [k]$, $\mathrm{Cl}(T_i|L(U)) = \{L(U(i))\} \cup \{L(v) : v$ is a descendant of a node in $U(i)\}$.*

Together with Lemma 1, Observation 2 shows that $T_i|L(U)$ is completely determined by the descendants of $U(i)$.

A valid subset U of $V(\mathcal{P})$ is *compatible* if there exists a phylogenetic tree T with $L(T) = L(U)$ that displays $T_i|L(U)$ for every $i \in [k]$. If such a tree T exists, we say that T *displays* U.

Lemma 3. *Profile $\mathcal{P}$ is compatible if and only if every valid subset of $V(\mathcal{P})$ is compatible.*

Proof. ($\Leftarrow$) If every valid subset of $V(\mathcal{P})$ is compatible, then, in particular, so is the set U_{init} of Equation (1). Let T be a tree that displays U_{init}. Then, $L(T) = L(U_{\text{init}}) = L(\mathcal{P})$. Thus, for every $i \in [k]$, $T_i|L(T) = T_i$, and thus T displays T_i. Hence, $\mathcal{P}$ is compatible.

($\Rightarrow$) Suppose $\mathcal{P}$ is compatible, but there is a valid subset U of $V(\mathcal{P})$ that is not compatible. Let T be a tree that displays $\mathcal{P}$. But then $T|U$ displays U, a contradiction. □

BUILDST (Algorithm 1), which is closely related to Semple and Steel's BUILD algorithm [16], determines whether a valid set $U \subseteq V(\mathcal{P})$ is compatible. The key difference between BUILDST and BUILD is that the latter uses the triplet graph $\Gamma(\mathcal{P})$, while BUILDST uses the graph $G_{\mathcal{P}}(U)$, for different subsets U of $V(\mathcal{P})$. As we show in Lemma 5, the two graphs are closely related. Nevertheless, $G_{\mathcal{P}}(U)$ offers some computational advantages over the triplet graph. Intuitively, this is because $G_{\mathcal{P}}(U)$ is a more compact representation of the triplets in $\mathcal{R}(\mathcal{P})$.

BUILDST(U) attempts to build a tree T_U for U. Step 1 initializes the root of T_U. If $L(U)$ consists of one or two species, then U is trivially compatible; Steps 2–5 handle these cases. The loop in lines 6–8 identifies the indices $i \in [k]$ such that $U(i)$ is a singleton. For each such i, it removes the single element v in $U(i)$ and replaces v by its children in T_i. As we argue in the proof of Theorem 1, when $\mathcal{P}$ is compatible, all such nodes v map to the same node w in the T that displays $\mathcal{P}$, in the sense that $L(w)$ is the smallest cluster in T such that $L(v) \subseteq L(w)$[2]. In Theorem 1, we also show that, if $G_{\mathcal{P}}(U)$ remains connected after steps 6–8, then U is incompatible. This case is handled in Line 11. Otherwise, Lines 12–17 recursively process each connected component of $G_{\mathcal{P}}(U)$. If the recursive calls succeed in finding trees for all the components, these trees are assembled into a phylogeny for U by joining them to the root created in Step 1. If any of the recursive calls determines that a component is incompatible, then U is declared to be incompatible.

The correctness of BUILDST relies on two lemmas, the first of which can be proved using induction.

Lemma 4. *If, given a valid set $U \subseteq V(\mathcal{P})$,* BUILDST$(U)$ *returns a tree T_U, then T_U is a phylogenetic tree such that $L(T_U) = L(U)$.*

The next lemma is central to the correctness proof of BUILDST.

Lemma 5. *Let $W_1, \ldots, W_p$ be the connected components of $G_{\mathcal{P}}(U)$ at step 9 of* BUILDST(U)*, for some valid set $U \subseteq V(\mathcal{P})$. Then,*

(i) for each $j \in [p]$, W_j is a valid set, and

(ii) the connected components of $\Gamma(\mathcal{P}|L(U))$ are precisely $L(W_1), \ldots, L(W_p)$.

Proof. (i) Let U_{bef} and U_{aft} denote the values of U before and after the executing steps 6–8. Each element of U_{aft} is either an element of U_{bef} or a child of some $v \in U_{\text{bef}}$. Indeed, in the latter case, every child of v is in U_{aft}. Thus, since, by assumption, U_{bef} is valid, and for every non-leaf node v, $L(v) = \bigcup_{w \in \text{Ch}(v)} L(w)$, U_{aft} must also be valid. Part (i) follows.

(ii) We first show that the following holds after steps 6–8.

[2]Thus, v plays the role of a *semi-universal node*, in the sense of Pe'er et al. [13].

Algorithm 1: BUILDST(U)

Input: A valid set $U \subseteq V(\mathcal{P})$.

Output: A tree T_U that displays U, if U is compatible; `incompatible` otherwise.

1 Create a node r_U
2 **if** $|L(U)| = 1$ **then**
3 | **return** the tree consisting of node r_U, labeled by the single species in $L(U)$
4 **if** $|L(U)| = 2$ **then**
5 | **return** the tree consisting of node r_U and two children, each labeled by a different species in $L(U)$
6 **foreach** $i \in [k]$ *such that* $|U(i)| = 1$ **do**
7 | Let v be the single element in $U(i)$
8 | $U = (U \setminus \{v\}) \cup \mathrm{Ch}(v)$
9 Let $W_1, W_2, \ldots, W_p$ be the connected components of $G_{\mathcal{P}}(U)$
10 **if** $p = 1$ **then**
11 | **return** `incompatible`
12 **foreach** $j \in [p]$ **do**
13 | Let $t_j = \text{BUILDST}(W_j)$
14 | **if** t_j *is a tree* **then**
15 | | Add t_j to the set of subtrees of r_U
16 | **else**
17 | | **return** `incompatible`
18 **return** the tree with root r_U

Claim 1. *Let a and b be any two species in $L(U)$. Then, (a, b) is an edge in $\Gamma(\mathcal{P}|L(U))$ if and only if there exists a node $v \in U$ such that $a, b \in L(v)$.*

Proof of claim. Observe that, after steps 6–8, $|U(i)| \neq 1$, for each $i \in [k]$.

($\Leftarrow$) Suppose that (a, b) is an edge in $\Gamma(\mathcal{P}|L(U))$. Then, there is an $i \in [k]$ such that $ab|x \in \mathcal{R}(T_i|L(U))$. Thus, there must be a proper descendant w of $r(T_i|L(U))$ such that $\{a, b\} \subseteq L(w)$. Observation 2 and the fact that $|U(i)| > 1$ imply that $L(w) \subseteq L(v)$ for some $v \in U(i)$.

($\Rightarrow$) Suppose that there is an $i \in [k]$ such that $a, b \in L(v)$ for some $v \in U(i)$. Choose a node $v' \in U(i) \setminus \{v\}$ — such a v' must exist, since $|U(i)| \geq 2$ — and choose some $x \in L(v')$. Then, $ab|x \in \mathcal{R}(T_i|L(U))$, and, hence, (a, b) is an edge of $\Gamma(\mathcal{P}|L(U))$. □

Observe that both $\Pi_1 = \{A : A \text{ is a connected component of } \Gamma(\mathcal{P}|L(U))\}$ and $\Pi_2 = \{L(W) : W \text{ is a connected component of } G_{\mathcal{P}}(U)\}$ are partitions of $L(U)$. We prove that $\Pi_1 = \Pi_2$ by showing that (a) for each connected component A of $\Gamma(\mathcal{P}|L(U))$ there exists a connected component W of $G_{\mathcal{P}}(U)$ such that $A \subseteq L(W)$, and (b) for each connected component W of $G_{\mathcal{P}}(U)$ there exists a connected component A of $\Gamma(\mathcal{P}|L(U))$ such that $L(W) \subseteq A$.

(a) Let A be any connected component of $\Gamma(\mathcal{P}|L(U))$. We argue that any two species a, b in A must be in the same connected component of $G_{\mathcal{P}}(U)$. Let $U_a = \{v \in U : a \in L(v)\}$ and $U_b = \{v \in U : b \in L(v)\}$. Then, each of U_a and U_b is a clique in $G_{\mathcal{P}}(U)$. It thus suffices to show that there is a path between some node in U_a and some node in U_b.

By the definition of A, there exists a path between a and b in $\Gamma(\mathcal{P}|L(U))$. Suppose this path is $\rho = \langle a_1, \ldots, a_m \rangle$, where $a_1 = a$ and $a_m = b$. By Claim 1, for each $l \in [m-1]$, there exists a node $w_l \in U$ such that $\{a_l, a_{l+1}\} \subseteq L(w_i)$. For each $l \in [m-2]$, $L(w_l) \cap L(w_{l+1}) \neq \emptyset$, so there is a edge between w_i and w_{i+1} in $G_{\mathcal{P}}(U)$. Hence, $\pi = \langle w_1, \ldots, w_{m-1} \rangle$ is a path from w_1 to w_{m-1} in $G_{\mathcal{P}}(U)$. By the definition of ρ, $a \in L(w_1)$ and $b \in L(w_{m-1})$, so $w_1 \in U_a$ and $w_l \in U_b$. This completes the proof of part (a).

(b) Let W be any connected component of $G_{\mathcal{P}}(U)$. If $|L(W)| = 1$, the statement holds trivially, so assume that $|L(W)| > 1$. We argue that any two species a, b in $L(W)$ are in the same connected component of $\Gamma(\mathcal{P}|L(U))$. Let v_a and v_b be nodes in W such that $a \in L(v_a)$ and $b \in L(v_b)$. If $v_a = v_b$, then, by Claim 1, (a, b) is an edge of $\Gamma(\mathcal{P}|L(U))$, and we are done. So, suppose instead that $v_a \neq v_b$.

Let us call a path π from v_a to v_b *good* if $|L(w)| > 1$ for every node w in π. We claim that there exists a good path from v_a to v_b. To prove this claim, we first argue that we can choose v_a and v_b such that $|L(v_a)|, |L(v_b)| > 1$. Indeed, consider the case of species a (the case for b is analogous). If $|L(v)| = 1$ for every node $v \in W$ such that $a \in L(v)$, then we would have $|L(W)| = 1$, contradicting our assumption that $|L(W)| > 1$. Now, suppose the path π from v_a to v_b has a node $w \notin \{v_a, v_b\}$ such that $|L(w)| = 1$. Let w' and w'' be the predecessor and successor of w in π. Then, $L(w') \cap L(w'') = L(w) \neq \emptyset$, so there is an edge between w' and w''. Thus, we can delete w from π and the resulting sequence remains a path between v_a and v_b.

Let $\pi = \langle w_1, \ldots, w_l \rangle$, where $w_1 = v_a$ and $w_l = v_b$, be a good path from v_a to v_b in $G_{\mathcal{P}}(U)$. Choose a sequence of species $\rho = \langle c_1, \ldots, c_{l+1} \rangle$, where $c_1 = a$, $c_{l+1} = b$ and, for each $j \in [l]$, $c_j, c_{j+1} \in L(w_j)$ and $c_j \neq c_{j+1}$. Note that such a choice is always possible. Then, by Claim 1, (c_j, c_{j+1}) is an edge of $\Gamma(\mathcal{P}|L(U))$. Hence, ρ is a path from a to b in $\Gamma(\mathcal{P}|L(U))$. □

We are now ready to prove the correctness of BUILDST.

Theorem 1. *Let U_{init} be the set defined in Equation* (1)*. Then,* BUILDST(U_{init}) *either (i) returns a tree T that displays $\mathcal{P}$, if $\mathcal{P}$ is compatible, or (ii) returns* `incompatible` *otherwise.*

Proof. We first argue that if BUILDST(U_{init}) outputs `incompatible`, $\mathcal{P}$ is indeed incompatible. Assume, on the contrary, that $\mathcal{P}$ is compatible. Then, there must be a call BUILDST(U) for some valid subset U such that $|L(U)| > 2$, in which the graph $G(U)$ of step 9 has a single connected component, $W_1 = U$. By Lemma 3, U must be compatible, so there exists a phylogeny T_U that displays U. By Observation 1, $\Gamma(\{T_U\})$ has at least two connected components A and B. By Lemma 5(ii), however, $\Gamma(\mathcal{P}|L(U))$ is connected, so there exist species $a \in A$ and $b \in B$ such that $ab|c \in \mathcal{R}(\mathcal{P}|U)$. But $ab|c \notin \mathcal{R}(T)$, and, by Lemma 2, T does not display some tree in $\mathcal{P}|L(U)$, a contradiction. Thus, $G(U)$ has at least two components.

Now, suppose that BUILDST(U_{init}) returns a tree T. We prove that T displays $\mathcal{P}$ by arguing that for each $i \in [k]$ there is a mapping $\phi_i : V(T_i) \to V(T)$ that maps every node $v \in V(T_i)$ to a node $\phi_i(v) \in V(T)$ such that $L(v) \subseteq L(\phi_i(v))$.

By Lemma 4, each recursive call BUILDST(U) returns a phylogenetic tree T_U for $L(U)$. Let r_U denote the root of T_U. We have two cases.

Case (i): $|L(U)| \leq 2$**.** For each $i \in [k]$, we must have $|U(i)| \in \{0, 1\}$. We only need to consider the case where $|U(i)| = 1$. Let v be the single node in $U(i)$. Note that $L(v) \subseteq L(r_U)$. Thus, we make $\phi_i(v) = r_U$. If $|L(U(i))| = 1$, we are done. Otherwise, $|L(U(i))| = 2$. Then, v has two children, v_1 and v_2, both leaves, labeled with, say, species s_1 and s_2, respectively. Node r_U also has two children, r_1 and r_2. Assume, without loss of generality, that these children are labeled with species s_1 and s_2, respectively. Then, $L(v_j) = L(r_j)$ for $j \in \{1, 2\}$. Therefore, we make $\phi_i(v_j) = r_j$ for each $j \in \{1, 2\}$.

Case (ii): $|L(U)| > 2$**.** Let U_{bef} be the value of U before entering the loop of lines 6–8, and let U_{aft} be the value of U at line 9, after the loop of lines 6–8 terminates. Let $U_{\text{rem}} = \{v \in U_{\text{bef}} : v \in U_{\text{bef}}(i)$ for

some $i \in [k]$ such that $|U_{\text{bef}}(i)| = 1\}$. Then $U_{\text{aft}} = (U_{\text{bef}} \setminus U_{\text{rem}}) \cup \{u \in \text{Ch}(v) : v \in U_{\text{rem}}\}$. Assume inductively that every descendant of a node in U_{aft} is mapped to an appropriate node in T_U. It therefore suffices to establish mappings for the nodes in U_{rem}. Now, for every $v \in U_{\text{rem}}$, $L(v) \subseteq L(r_U)$. Thus, we make $\phi(v) = r_U$ for every $v \in U_{\text{rem}}$. □

4 Implementation

We now explain how to implement BUILDST in order to solve the tree compatibility problem in $O(M_{\mathcal{P}} \log^2 M_{\mathcal{P}})$ time. Consider a call to BUILDST(U). Recall that we can assume that $G_{\mathcal{P}}(U)$ is connected. BUILDST(U) requires the following three pieces of information.

(G1) *The value of* $|L(U)|$. This number is needed in Lines 2 and 4 of BUILDST.

(G2) *The set* $J(U)$ *of all* $i \in [k]$ *such that* $|U(i)| = 1$. Set $J(U)$ contains the indices i considered in Lines 6–8 of BUILDST.

(G3) *The set* $U(i) = U \cap V(T_i)$ *for each* $i \in [k]$. For each $i \in J(U)$, $U(i)$ contains precisely the element v used in Lines 7 and 8 of BUILDST.

It is straightforward to obtain (G1), (G2), and (G3) for the valid set U_{init} of Equation (1): $|L(U_{\text{init}})| = n$, $J(U_{\text{init}}) = [k]$, and, for every $i \in [k]$, $U_{\text{init}}(i) = \{r(T_i)\}$. Now assume that we have (G1), (G2), and (G3) at the beginning of some call to BUILDST(U). Steps 6–8 modify U and, therefore, $G_{\mathcal{P}}(U)$. Suppose that, after Line 9, $G_{\mathcal{P}}(U)$ has more than one connected component. We need to compute (G1), (G2), and (G3) for each connected component, in order to pass this information to the recursive calls in Line 13. That is, if $p > 1$, for each $j \in [p]$, we need to compute $|L(W_j)|$, $J(W_j)$, and $W_j(i) = W_j \cap V(T_i)$, for each $i \in [k]$.

We use the dynamic graph connectivity data structure by Holm et al. [11]. We refer to this data structure as *HDT*. HDT guarantees that, if we start with no edges in a graph with N vertices, the amortized cost of each update is $O(\log^2 N)$. For efficiency, however, do not use HDT directly on $G_{\mathcal{P}}(U)$. The reason is that the edges of $G_{\mathcal{P}}(U)$ are defined via intersections of sets of species, which could make it costly to determine the new nodes and edges created as a result of Step 8. We avoid this problem through an indirect approach that uses an auxiliary graph $H_{\mathcal{P}}(U)$, defined below. As we shall see, $H_{\mathcal{P}}(U)$ offers another advantage over $G_{\mathcal{P}}(U)$: maintaining $H_{\mathcal{P}}(U)$ only requires handling deletions, but maintaining $G_{\mathcal{P}}(U)$ additionally requires handling insertions.

We define $H_{\mathcal{P}}(U)$ as a subgraph of the graph $H_{\mathcal{P}}$ constructed as follows. For each species $s \in L(\mathcal{P})$, create a new node $x_s \notin V(\mathcal{P})$, and let $X_{\mathcal{P}} = \{x_s : s \in L(\mathcal{P})\}$. Then, $H_{\mathcal{P}}$ is the graph whose vertex set is $V(\mathcal{P}) \cup X_{\mathcal{P}}$ and whose edge set is $E(\mathcal{P}) \cup \{(u, x_s) : u$ is a leaf in T_i, for some $i \in [k]$, such that $\lambda(u) = s\}$. Note that $H_{\mathcal{P}}$ has $O(M_{\mathcal{P}})$ nodes and edges, and can be constructed from $\mathcal{P}$ in $O(M_{\mathcal{P}})$ time. $H_{\mathcal{P}}$ is essentially the *display graph* for $\mathcal{P}$ [4]. The display graph is the result of glueing together leaves in $\mathcal{P}$ labeled by the same species. Contrast this with $H_{\mathcal{P}}$, which connects leaves with a common label through nodes in $X_{\mathcal{P}}$. This minor difference with respect to the display graph serves to simplify our presentation.

Given a valid subset U of $V(\mathcal{P})$, we define $H_{\mathcal{P}}(U)$ as the subgraph of $H_{\mathcal{P}}$ induced by $\{v : v$ is a descendant of some node $u \in U\} \cup \{x_s \in X_{\mathcal{P}} : s \in L(U)\}$. Note that $H_{\mathcal{P}}(U_{\text{init}}) = H_{\mathcal{P}}$. See Figure 2.

The next result states the basic properties of $H_{\mathcal{P}}(U)$.

Lemma 6. *The following statements hold for any valid subset* U *of* $V(\mathcal{P})$.

(i) Let v *be a node in* U. *If* $U' = (U \setminus \{v\}) \cup \text{Ch}(v)$, *then* $H_{\mathcal{P}}(U')$ *is obtained from* $H_{\mathcal{P}}(U)$ *by deleting* v *and every edge* (v, u) *such that* $u \in \text{Ch}(v)$.

(ii) Any two nodes in U *are in the same connected component in* $G_{\mathcal{P}}(U)$ *if and only if they are in the same connected component of* $H_{\mathcal{P}}(U)$.

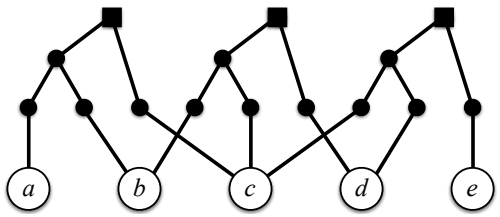

Figure 2: The graph $H_{\mathcal{P}}(U_{\text{init}})$ for the profile $\mathcal{P}$ of Figure 1. The nodes of U_{init} are drawn as squares. Nodes in the set $\{x_s : s \in L(\mathcal{P})\}$ are labeled with the corresponding species. Species labeling the leaves of trees in $\mathcal{P}$ are omitted.

Proof. The proof of part (i) is trivial, so we focus on part (ii). We argue that, for any two nodes $v, w \in U$, v and w are in the same connected component in $G_{\mathcal{P}}(U)$ if and only if they are in the same connected component of $H_{\mathcal{P}}(U)$.

($\Leftarrow$) Suppose v and w are in the same connected component of $H_{\mathcal{P}}(U)$. Then, there exists a path π between v and w in $H_{\mathcal{P}}(U)$. Let $\pi' = \langle u_1, \ldots, u_m \rangle$ be the sequence obtained by striking out from π all nodes not in U. Note that $u_1 = v$ and $u_m = w$. For each $i \in [m-1]$, the subpath of π between u_i and u_{i+1} contains at least one node from $X_{\mathcal{P}}$, say x_s, for some $s \in L(\mathcal{P})$. But then $s \in L(u_i) \cap L(u_{i+1})$, so there is an edge between u_i and u_{i+1} in $G_{\mathcal{P}}(U)$. Thus, π' is a path between u_i and u_{i+1} in $G_{\mathcal{P}}(U)$. Hence, v and w are in the same connected component in $G_{\mathcal{P}}(U)$.

($\Rightarrow$) Suppose v and w are in the same connected component of $G_{\mathcal{P}}(U)$. Then, there exists a path between v and w in $G_{\mathcal{P}}(U)$. Let $\langle u_1, \ldots, u_m \rangle$ be one such path, where $u_1 = v$ and $u_m = w$. By definition of $G_{\mathcal{P}}(U)$, for each $i \in [m-1]$, $L(u_i) \cap L(u_{i+1}) \neq \emptyset$. Pick any species $s \in L(u_i) \cap L(u_{i+1})$. Hence, in $H_{\mathcal{P}}(U)$ there is a path from u_i to u_{i+1} that goes through x_s. Since this holds for each $i \in [m-1]$, there exists a path from $u_1 = v$ to $u_m = w$ in $H_{\mathcal{P}}(U)$. Thus, v and w are in the same connected component in $H_{\mathcal{P}}(U)$. □

By Lemma 6(ii), the connected components $W_1, \ldots, W_p$ of $G_{\mathcal{P}}(U)$ can be put into a one-to-one correspondence with the connected components $Y_1, \ldots, Y_p$ of $H_{\mathcal{P}}(U)$ so that $W_j = Y_j \cap U$ for each $j \in [p]$.

We represent $H_{\mathcal{P}}(U)$ using the aforementioned HDT data structure. For each connected component Y of $H_{\mathcal{P}}(U)$, we maintain three fields:

(H1) Y.`count`, the cardinality of $Y \cap X_{\mathcal{P}}$,

(H2) Y.`singleton`, a doubly-linked list that contains all indices $i \in [k]$ such that $|U(i)| = 1$, and

(H3) Y.`List`, an array where, for each $i \in [k]$, Y.`List`$[i]$ is a doubly-linked list consisting of the elements of $Y \cap U(i)$.

Recall that we assume that $G_{\mathcal{P}}(U)$ is connected at the beginning of a call to BUILDST(U). Thus, by Lemma 6, $H_{\mathcal{P}}(U)$ has a single connected component, Y. Then, $|L(U)| = Y$.`count`, $J(U) = Y$.`singleton`, and Y.`List`$[i]$ contains the elements of $U(i)$, for each $i \in [k]$. Thus, the three fields of Y provide BUILDST(U) with the information that it needs — that is, (G1), (G2), and (G3). In particular, they allow us to easily find each node v considered in Line 7 of BUILDST(U). Line 8 is then performed as a series of edge deletions, one for each edge (v, u) such that $u \in \text{Ch}(v)$, followed by the deletion of v (we provide further details below). By Lemma 6(i), this correctly updates $H_{\mathcal{P}}(U)$. The deletions break $H_{\mathcal{P}}(U)$ down into a collection of connected components $Y_1, \ldots, Y_p$. For each $j \in [p]$, Y_j corresponds to a connected component W_j of $G_{\mathcal{P}}(U)$ that (if $p > 1$) is processed in a recursive call in Line 13. We need to compute Y_j.`count`, Y_j.`singleton`, Y_j.`List` for each $j \in [p]$, in order to provide this information to the recursive calls.

The total number of edge and node deletions executed by BUILDST(U_{init}) — including all deletions conducted by the recursive calls — cannot exceed the total number of edges and nodes in $H_{\mathcal{P}}$, which is $O(M_{\mathcal{P}})$. The HDT data structure allows us to maintain connectivity information throughout the entire algorithm in $O(M_{\mathcal{P}} \log^2 M_{\mathcal{P}})$. In the remainder of this section, we show that we can maintain the `count`, `singleton`, and `List` fields throughout the entire algorithm in total time $O(M_{\mathcal{P}} \log^2 M_{\mathcal{P}})$. We also argue that all the required information for $H_{\mathcal{P}}(U_{\text{init}})$ can be initialized in $O(M_{\mathcal{P}})$ time.

Let $Y_{\text{init}} = V(\mathcal{P}) \cup X_{\mathcal{P}}$ be the vertex set of $H_{\mathcal{P}}(U_{\text{init}})$. Then, Y_{init} is the single connected component of $H_{\mathcal{P}}(U_{\text{init}})$. We initialize the data fields of Y_{init} as follows: (1) $Y_{\text{init}}.\texttt{count} = |L(\mathcal{P})|$, (2) $Y_{\text{init}}.\texttt{singleton}$ is the set $[k]$, and (3) for each $i \in [k]$, $Y_{\text{init}}.\texttt{List}[i]$ consists of $r(T_i)$. Thus, we can initialize all data fields in $O(M_{\mathcal{P}})$ time.

We assume that every node v in $H_{\mathcal{P}}(U)$ is either *marked*, if $v \in U$, or *unmarked*, if $v \notin U$. Initially, each node $v \in U_{\text{init}}$ is marked, and every node $v \in Y_{\text{init}} \setminus U_{\text{init}}$ is unmarked. We also assume that for each node v in $H_{\mathcal{P}}(U)$, we maintain sufficient information to be able to determine in $O(1)$ time whether $v \in X_{\mathcal{P}}$ or $v \in V(\mathcal{P})$, and that, in the latter case, we have $O(1)$-time access to the index $i \in [k]$ such that $v \in V(T_i)$. For each i such that $Y.\texttt{List}[i]$ contains exactly one element, we maintain a pointer from $Y.\texttt{List}[i]$ to the entry for i in $Y.\texttt{singleton}$. This allows us to update $Y.\texttt{singleton}$ in $O(1)$ time when $U(i)$ is no longer a singleton. For each marked node $v \in Y$ (so $v \in U$), we maintain a pointer from v to the element in $Y.\texttt{List}[i]$ that contains v. This allows us to update $Y.\texttt{List}[i]$ in $O(1)$ time when v becomes unmarked.

Consider a call to BUILDST(U) for some valid set U. Step 1 takes $O(1)$ time. Since $H_{\mathcal{P}}(U)$ initially consists of a single connected component, say Y, and we have $Y.\texttt{count}$, Steps 2–5 also take $O(1)$ time. Let $H = H_{\mathcal{P}}(U)$. We implement the loop in lines 6–8 as follows. First, we enumerate the indices in $J = J(U)$ in $O(|J|)$ time by listing the elements of $Y.\texttt{singleton}$. For each $i \in J$, we retrieve and remove the single element v_i of $U(i)$ from $Y.\texttt{List}[i]$, and then delete i from $Y.\texttt{singleton}$. This takes $O(1)$ time. We unmark v_i, and for every node $u \in \text{Ch}(v_i)$ we mark u and add it to $Y.\texttt{List}[i]$. This takes $O(1)$ time per edge. We then successively delete each edge (v_i, u) such that $u \in \text{Ch}(v_i)$, updating (H1)–(H3) for each newly-created component along the way. Once these edges are deleted, we delete v_i itself. By Lemma 6(i), the result is the graph $H_{\mathcal{P}}(U)$ for the new set U. Let us focus on how to handle the deletion of a single edge $e = (v_i, u)$.

Let Y' be the connected component of H that currently contains v_i. We query the HDT data structure to determine, in $O(\log^2 M_{\mathcal{P}})$ amortized time, whether deleting (v_i, u) splits Y' into two components. If Y' remains connected, no updates are needed. Otherwise, Y' is split into two parts Y_1 and Y_2. To fill in the `count`, `singleton`, and `List` fields of Y_1 and Y_2, we use the well-known technique of scanning the smaller component [7]. The HDT data structure maintains the sizes of the various components [11], so we can determine in $O(1)$ which of Y_1 and Y_2 has fewer nodes. Suppose without loss of generality that $|Y_1| \leq |Y_2|$. We initialize $Y_2.\texttt{count}$ and $Y_2.\texttt{List}$ to $Y'.\texttt{count}$ and $Y'.\texttt{List}$, respectively. We initialize $Y_1.\texttt{count}$ to 0 and $Y_1.\texttt{List}[i]$ to null for each $i \in [k]$. We then scan each node v in Y_1, and do the following. If $v \in X_{\mathcal{P}}$, we decrement $Y_2.\texttt{count}$ and increment $Y_1.\texttt{count}$. Otherwise $v \in V(\mathcal{P})$; assume that $v \in V(T_i)$. If v is marked, we remove v from $Y_2.\texttt{List}[i]$ and add v to $Y_1.\texttt{List}[i]$. This operation requires at most one update in each of $Y_1.\texttt{singleton}$ and $Y_2.\texttt{singleton}$; each update takes $O(1)$ time.

We claim that any node v is scanned $O(\log M_{\mathcal{P}})$ times over the entire execution of BUILDST(U_{init}). To verify this, let $N(v)$ be the number of nodes in the connected component containing v. Suppose that, initially, $N(v) = N$. Then, the rth time we scan v, $N(v) \leq N/2^r$. Thus, v is scanned $O(\log N)$ times. The claim follows, since $N = O(M_{\mathcal{P}})$. Therefore, the total number of updates over all nodes is $O(M_{\mathcal{P}} \log M_{\mathcal{P}})$, and the work per update is $O(1)$.

To summarize, the work done by BUILDST consists of three parts: (i) initialization, (ii) maintaining connected components, and (iii) maintaining the `count`, `singleton`, and `List` fields for each connected

component. Part (i) takes $O(M_{\mathcal{P}})$ time. Part (ii) involves $O(M_{\mathcal{P}})$ edge and node deletions on the HDT data structure, at an amortized cost of $O(\log^2 M_{\mathcal{P}})$ per deletion. Part (iii) involves scanning the nodes of our graph every time a deletion creates a new component, for a total of $O(M_{\mathcal{P}} \log M_{\mathcal{P}})$ scans, at $O(1)$ cost per scan, over the entire execution of BUILDST. This yields our main result.

Theorem 2. *Let U_{init} be the set defined in Equation* (1)*. Then, there exists and implementation of* BUILDST *such that* BUILDST(U_{init}) *runs in $O(M_{\mathcal{P}} \log^2 M_{\mathcal{P}})$ time.*

5 Discussion

A trivial lower bound for the tree compatibility problem is $\Omega(M_{\mathcal{P}})$, the time to read the input. Thus, our result leaves us a polylogarithmic factor away from an optimal algorithm for compatibility. Is it possible to reduce or even eliminate this gap? The bottleneck is the time to maintain the information associated with the various components of $H_{\mathcal{P}}(U)$. It is conceivable that the special structure of this graph and the way the deletions are performed could be used to our advantage. A second question is how well our algorithm performs in practice. To investigate this, it should be possible to leverage existing knowledge on the empirical behavior of dynamic connectivity data structures [12].

References

[1] A. V. Aho, Y. Sagiv, T. G. Szymanski, and J. D. Ullman. Inferring a tree from lowest common ancestors with an application to the optimization of relational expressions. *SIAM J. Comput.*, 10(3):405–421, 1981.

[2] B. R. Baum. Combining trees as a way of combining data sets for phylogenetic inference, and the desirability of combining gene trees. *Taxon*, 41:3–10, 1992.

[3] O. R. P. Bininda-Emonds, M. Cardillo, K. E. Jones, R. D. E. MacPhee, R. M. D. Beck, R. Grenyer, S. A. Price, R. A. Vos, J. L. Gittleman, and A. Purvis. The delayed rise of present-day mammals. *Nature*, 446:507–512, 2007.

[4] D. Bryant and J. Lagergren. Compatibility of unrooted phylogenetic trees is FPT. *Theoretical Computer Science*, 351:296–302, 2006.

[5] P. Buneman. A characterisation of rigid circuit graphs. *Discrete Math.*, 9:205–212, 1974.

[6] B. Courcelle. The monadic second-order logic of graphs I. Recognizable sets of finite graphs. *Inf. Comput.*, 85(1):12–75, 1990.

[7] S. Even and Y. Shiloach. An on-line edge-deletion problem. *J. ACM*, 28(1):1–4, Jan. 1981.

[8] S. Grünewald, M. Steel, and M. S. Swenson. Closure operations in phylogenetics. *Mathematical Biosciences*, 208:521—537, 2007.

[9] M. R. Henzinger, V. King, and T. Warnow. Constructing a tree from homeomorphic subtrees, with applications to computational evolutionary biology. *Algorithmica*, 24:1–13, 1999.

[10] C. E. Hinchliff, S. A. Smith, J. F. Allman, J. G. Burleigh, R. Chaudhary, L. M. Coghill, K. A. Crandall, J. Deng, B. T. Drew, R. Gazis, K. Gude, D. S. Hibbett, L. A. Katz, H. D. Laughinghouse IV, E. J. McTavish, P. E. Midford, C. L. Owen, R. H. Reed, J. A. Reesk, D. E. Soltis, T. Williams, and K. A. Cranston. Synthesis of phylogeny and taxonomy into a comprehensive tree of life. *Proceedings of the National Academy of Sciences*, 2015. In press.

[11] J. Holm, K. de Lichtenberg, and M. Thorup. Poly-logarithmic deterministic fully-dynamic algorithms for connectivity, minimum spanning tree, 2-edge, and biconnectivity. *J. ACM*, 48(4):723–760, July 2001.

[12] R. Iyer, D. Karger, H. Rahul, and M. Thorup. An experimental study of polylogarithmic, fully dynamic, connectivity algorithms. *J. Exp. Algorithmics*, 6, Dec. 2001.

[13] I. Pe'er, T. Pupko, R. Shamir, and R. Sharan. Incomplete directed perfect phylogeny. *SIAM J. Comput.*, 33(3):590–607, 2004.

[14] M. A. Ragan. Phylogenetic inference based on matrix representation of trees. *Molecular Phylogenetics and Evolution*, 1:53–58, 1992.

[15] M. J. Sanderson. Phylogenetic signal in the eukaryotic tree of life. *Science*, 321(5885):121–123, 2008.

[16] C. Semple and M. Steel. *Phylogenetics*. Oxford Lecture Series in Mathematics. Oxford University Press, Oxford, 2003.

[17] M. A. Steel. The complexity of reconstructing trees from qualitative characters and subtrees. *J. Classification*, 9:91–116, 1992.